

A family of statistical symmetric divergences based on Jensen's inequality

Frank Nielsen ^{a,b,*}

^a*Ecole Polytechnique
Computer Science Department (LIX)
Palaiseau, France.*

^b*Sony Computer Science Laboratories, Inc. (FRL),
3-14-13 Higashi Gotanda 3F, Shinagawa-Ku,
Tokyo 141-0022, Japan.*

Abstract

We introduce a novel parametric family of symmetric information-theoretic distances based on Jensen's inequality for a convex functional generator. In particular, this family unifies the celebrated Jeffreys divergence with the Jensen-Shannon divergence when the Shannon entropy generator is chosen. We then design a generic algorithm to compute the unique centroid defined as the minimum average divergence. This yields a smooth family of centroids linking the Jeffreys to the Jensen-Shannon centroid. Finally, we report on our experimental results.

Key words: Kullback-Leibler divergence ; Jensen-Shannon divergence ; centroid

* Corresponding author. Fax # (+81) 3-5448-4380.

Email address: Frank.Nielsen@acm.org (Frank Nielsen).

URL: <http://www.sonycs1.co.jp/person/nielsen/> (Frank Nielsen).

1 Introduction to statistical distances

The Shannon differential entropy [Cover and Thomas(1991)] of a continuous probability distribution¹ p measures the amount of uncertainty:

$$H(p) = \int p(x) \log \frac{1}{p(x)} dx = - \int p(x) \log p(x) dx. \quad (1)$$

The cross-entropy [Cover and Thomas(1991)] measures the amount of extra bits required to compute a code based on an observed empirical probability \tilde{p} instead of the true probability p (hidden by nature):

$$H(p : \tilde{p}) = \int p(x) \log \frac{1}{\tilde{p}(x)} dx = - \int p(x) \log \tilde{p}(x) dx. \quad (2)$$

The “:” notation emphasizes on the *oriented* aspect [Cover and Thomas(1991)] of the functional: $H(p : q) \neq H(q : p)$. The Kullback-Leibler divergence [Kullback and Leibler(1951), Cover and Thomas(1991)] is a *statistical distance measure* computing the *relative entropy* as follows:

$$\text{KL}(p : q) = \int p(x) \log \frac{p(x)}{q(x)} dx \quad (3)$$

$$= H(p : q) - H(p) \geq 0, \quad (4)$$

This last inequality is called Gibb’s inequality [Cover and Thomas(1991)], with equality if and only if $p = q$. We have $H(p : q) = H(p) + \text{KL}(p : q)$. The Kullback-Leibler divergence can be extended to unnormalized positive distributions (or positive arrays in discrete cases) as follows:

$$\text{eKL}(p : q) = \int \left(p(x) \log \frac{p(x)}{q(x)} + q(x) - p(x) \right) dx, \quad (5)$$

$$= \text{eH}(p : q) - \text{eH}(p) \geq 0, \quad (6)$$

with $\text{eH}(p : q) = \int (p(x) \log \frac{1}{q(x)} + q(x)) dx$ and $\text{eH}(p) = \text{eH}(p, p)$.

(Rényi based on an axiomatic approach [Rényi(1961)] derived yet another expression for the Kullback-Leibler divergence of unnormalized generalized distributions.)

¹ For sake of simplicity and without loss of generality, we consider the probability density function p of a continuous random variable $X \sim p$. For multivariate densities p , the integral notation \int denote the corresponding multi-dimensional integral, so that we write for short $\int p(x) dx = 1$. Our results hold for probability mass functions, and probability measures in general.

Many applications in Information Retrieval [Deselaers, Keysers, and Ney(2008), Rubner, Puzicha, Tomasi (IR)] requires to deal with a *symmetric* distortion measure. Jeffreys divergence [Jeffreys(1973)] (also called *J-divergence*) symmetrizes the oriented Kullback-Leibler divergence as follows:

$$J(p, q) = \text{KL}(p : q) + \text{KL}(q : p) = J(q, p) \quad (7)$$

$$= H(p : q) + H(q : p) - (H(p) + H(q)), \quad (8)$$

$$= \int (p(x) - q(x)) \log \frac{p(x)}{q(x)} dx. \quad (9)$$

Here, we replaced “.” by “,” in the distortion measure to emphasize the symmetric property: $J(p, q) = J(q, p)$. Jeffreys divergence is interpreted as *twice the average of the cross-entropies minus the average of the entropies*. One of the drawbacks of Jeffreys divergence is that it may be *unbounded* and therefore numerically quite unstable to compute in practice: For example, let $p = (p^i)_{i=1}^d$ and $q = (q^i)_{i=1}^d$ be frequency histograms with d bins (discrete distributions called multinomials), then $J(p, q) \rightarrow \infty$ if there exists one bin $l \in \{1, \dots, d\}$ such that p^l is above some constant, and $q^l \rightarrow 0$. In that case, $p^l \log \frac{p^l}{q^l} \rightarrow \infty$. To circumvent this unboundedness problem, the Jensen-Shannon divergence was introduced in [Lin(1991)]. The Jensen-Shannon divergence symmetrizes the Kullback-Leibler divergence by taking the *average relative entropy of the source distributions to the entropy of the average distribution* $\frac{p+q}{2}$:

$$\text{JS}(p, q) = \frac{1}{2} \left(\text{KL} \left(p : \frac{p+q}{2} \right) + \text{KL} \left(q : \frac{p+q}{2} \right) \right) = \text{JS}(q, p) \quad (10)$$

$$= \frac{1}{2} \left(H \left(p : \frac{p+q}{2} \right) - H(p) + H \left(q : \frac{p+q}{2} \right) - H(q) \right), \quad (11)$$

$$= \frac{1}{2} \int \left(p(x) \log \frac{2p(x)}{p(x) + q(x)} + q(x) \log \frac{2q(x)}{p(x) + q(x)} \right) dx, \quad (12)$$

$$= H \left(\frac{p+q}{2} \right) - \frac{H(p) + H(q)}{2} \geq 0. \quad (13)$$

The Jensen-Shannon divergence has always finite values, and its square root yields a metric, satisfying the triangular inequality. Moreover, we have the following information-theoretic inequality [Lin(1991)]

$$0 \leq \text{JS}(p, q) \leq \frac{1}{4} J(p, q). \quad (14)$$

By introducing the *K-divergence* [Lin(1991)] (see Eq. 7):

$$K(p : q) = \int p(x) \log \frac{2p(x)}{p(x) + q(x)} dx = \text{KL} \left(p : \frac{p+q}{2} \right), \quad (15)$$

we interpret the Jensen-Shannon divergence as the Jeffreys symmetrization of the K -divergence (see Eq. 7).

$$\text{JS}(p, q) = \frac{1}{2}(K(p : q) + K(q : p)), \quad (16)$$

$$= H\left(\frac{p+q}{2}\right) - \frac{H(p) + H(q)}{2}. \quad (17)$$

Consider the *skewed* K -divergence

$$K_\alpha(p : q) = \int p(x) \log \frac{p(x)}{(1-\alpha)p(x) + \alpha q(x)} dx, \quad (18)$$

and its symmetrized divergence

$$\text{JS}_\alpha(p, q) = \frac{K_\alpha(p : q) + K_\alpha(q : p)}{2} = \text{JS}_\alpha(q, p). \quad (19)$$

For $\alpha = \frac{1}{2}$, we find the Jensen-Shannon divergence: $\text{JS}(p, q) = \text{JS}_{\frac{1}{2}}(p, q)$. For $\alpha = 1$, we obtain half of Jeffreys divergence: $\text{JS}_1(p, q) = \frac{1}{2}J(p, q)$. It turns out that this family of α -Jensen-Shannon divergence belongs to a broader family of information-theoretic measures, called Ali-Silvey-Csiszár divergences [Csiszár(1967), Ali and Silvey(1966)]. A ϕ -divergence is defined for a strictly convex function ϕ such that $\phi(1) = 0$ as:

$$I_\phi(p : q) = \int q(x) \phi\left(\frac{p(x)}{q(x)}\right) dx. \quad (20)$$

We can always symmetrize ϕ -divergences by taking the *coupled* convex function $\phi^*(x) = x\phi(\frac{1}{x})$. Indeed, we get

$$I_{\phi^*}(p : q) = \int q(x) \phi^*\left(\frac{p(x)}{q(x)}\right) dx, \quad (21)$$

$$= \int q(x) \frac{p(x)}{q(x)} \phi\left(\frac{q(x)}{p(x)}\right) dx, \quad (22)$$

$$= \int p(x) \phi\left(\frac{q(x)}{p(x)}\right) dx = I_\phi(q : p). \quad (23)$$

Therefore, $I_{\phi+\phi^*}(p, q)$ is a symmetric divergence. Let $\phi^s = \phi + \phi^*$ denote the symmetrized generator. Jeffreys divergence is a ϕ -divergence for $\phi(u) = -\log u$ (and $\phi^s(u) = (u-1)\log u$). Similarly, Jensen-Shannon divergence is interpreted as $\text{JS}(p, q) = \frac{1}{2}(K(p : q) + K(q : p))$, with $\frac{1}{2}K(p : q)$ a ϕ -divergence

for $\phi(u) = \frac{u}{2} \log \frac{2u}{1+u}$, see [Lin(1991)]. It follows that Jensen-Shannon is also a ϕ -divergence. The α -Jensen-Shannon divergences are ϕ -divergences for the generators $\phi_\alpha^s = \phi_\alpha^* + \phi_\alpha$, with $\phi_\alpha^*(x) = -\log((1-\alpha) + \alpha x)$ and $\phi_\alpha(x) = -x \log((1-\alpha) + \frac{\alpha}{x})$. α -Jensen-Shannon divergences are convex statistical distances in both arguments.

One drawback for estimating α -JS divergences on *continuous* parametric densities (say, Gaussian distributions), is that the mixture of two Gaussians is not anymore a Gaussian, and therefore the average distribution falls outside the family of considered distributions. This explains the lack of closed-form solution for computing the Jensen-Shannon divergence on Gaussians.

Next, we introduce a novel family of symmetrized divergences which yields closed-form formulas for statistical distances of a large class of parametric distributions, called statistical exponential families.

2 A novel parametric family of Jensen divergences

At the heart of many statistical distances lies the celebrated Jensen's convex inequality [Jensen(1906)]. For a strictly convex function F and a parameter $\alpha \in \mathbb{R} \setminus \{0, 1\}$, let us define the α -skew Jensen divergence as

$$J_F^{(\alpha)}(p : q) = \frac{1}{\alpha(1-\alpha)} \int ((1-\alpha)F(p(x)) + \alpha F(q(x)) - F((1-\alpha)p(x) + \alpha q(x))) dx. \quad (24)$$

This statistical divergence is said *separable* as it can be rewritten as

$$J_F^{(\alpha)}(p : q) = \int j_F^{(\alpha)}(p(x) : q(x)) dx, \quad (25)$$

with the scalar basic distance being defined as

$$j_F^{(\alpha)}(x : y) = \frac{1}{\alpha(1-\alpha)} ((1-\alpha)F(x) + \alpha F(y) - F((1-\alpha)x + \alpha y)). \quad (26)$$

Furthermore, we refine the definition of α -skew Jensen divergence to real-valued d -dimensional vectors p and q as

$$J_F^{(\alpha)}(p : q) = \sum_{i=1}^d j_F^{(\alpha)}(p^i : q^i) dx. \quad (27)$$

In the limit cases, we find the oriented Kullback-Leibler divergences [Nielsen and Boltz(2011)] when we choose generator $F(x) = -x \log x$:

$$\lim_{\alpha \rightarrow 0} J_F^{(\alpha)}(p : q) = \text{KL}(p : q), \quad (28)$$

$$\lim_{\alpha \rightarrow 1} J_F^{(\alpha)}(p : q) = \text{KL}(q : p). \quad (29)$$

Observe also that $J_F^{(\alpha)}(q : p) = J_F^{(1-\alpha)}(p : q)$, and that therefore α -skew Jensen divergences are asymmetric distortion measures (except for $\alpha = \frac{1}{2}$). Therefore, let us symmetrize those α -skew divergences by averaging the two orientations as follows:

$$\text{sj}_F^{(\alpha)}(p, q) = \frac{1}{2}(J_F^{(\alpha)}(p : q) + J_F^{(\alpha)}(q : p)) \quad (30)$$

$$= \frac{1}{2}(J_F^{(\alpha)}(p : q) + J_F^{(1-\alpha)}(p : q)) \quad (31)$$

$$= \frac{1}{2\alpha(1-\alpha)} \int (F(p(x)) + F(q(x)) - F(\alpha p(x) + (1-\alpha)q(x)) - F((1-\alpha)p(x) + \alpha q(x))) dx \quad (32)$$

$$= \text{sj}_F^{(\alpha)}(q, p) = \text{sj}_F^{(1-\alpha)}(p, q) \geq 0 \quad (33)$$

For discrete d -dimensional parameter vectors p and q (with respective coordinates p^1, \dots, p^d and q^1, \dots, q^d), we define analogously the α -skew divergences as follows:

$$\text{sj}_F^{(\alpha)}(p, q) = \frac{1}{2\alpha(1-\alpha)} \sum_{i=1}^d (F(p^i) + F(q^i) - F(\alpha p^i + (1-\alpha)q^i) - F((1-\alpha)p^i + \alpha q^i)) dx. \quad (34)$$

Those statistical divergences are *separable* as they can be written as $\text{sj}_F^{(\alpha)}(p, q) = \int \text{sj}_F^{(\alpha)}(p(x) : q(x)) dx$, where sj_F denotes the corresponding basic scalar distance measure:

$$\text{sj}_F^{(\alpha)}(x, y) = \frac{1}{2\alpha(1-\alpha)} (F(x) + F(y) - F(\alpha x + (1-\alpha)y) - F((1-\alpha)x + \alpha y)) dx. \quad (35)$$

Figure 1 shows graphically this novel family of symmetric Jensen divergences by depicting its associated basic scalar distance $\text{sj}_F^{(\alpha)}$ (it is enough to consider $\alpha \in [0, \frac{1}{2}]$). Note that except for $\alpha \in \{0, 1\}$, this family of divergences have necessarily the boundedness property: $0 \leq \text{sj}_F^{(\alpha)}(p, q) < \infty, \forall \alpha \notin \{0, 1\}$

Consider the strict convex generator $F(x) = x \log x$ (Shannon information function or equivalently the negative Shannon entropy functional). Rewriting the divergence for this generator, we get a family of *symmetric Kullback-Leibler*




Fig. 1. A family of separable symmetric scalar Jensen divergences $\{sJ_F^{(\alpha)}(p, q) = \int sj_F^{(\alpha)}(p(x), q(x))dx\}_\alpha$ for $\alpha \in (0, \frac{1}{2}]$ based on Jensen's convexity gap that includes both Jeffreys divergence in the limit case $\alpha = 0$ and Jensen-Shannon divergence for $\alpha = \frac{1}{2}$, for the Shannon information generator. Here, we plot the scalar base distance $sj_F^{(\alpha)}$ that induces the statistical distance.

divergences:

$$sKL^{(\alpha)}(p, q) = \frac{1}{2\alpha(1-\alpha)} (H(\alpha p + (1-\alpha)q) + H((1-\alpha)p + \alpha q) - (H(p) + H(q))) \geq 0 \quad (36)$$

We have in the limit case:

$$\lim_{\alpha \rightarrow 0} sKL^{(\alpha)}(p, q) = J(p, q) = sKL^{(0)}(p, q). \quad (37)$$

That is, symmetrized α -Jensen divergences tend asymptotically to the Jeffreys divergence for the Shannon information generator. Furthermore, consider the case $\alpha = \frac{1}{2}$:

$$sKL^{(\frac{1}{2})}(p, q) = 2 \left(2H\left(\frac{p+q}{2}\right) - (H(p) + H(q)) \right) = 4JS(p, q). \quad (38)$$

Thus this family of symmetric Kullback-Leibler divergences unify both Jensen-Shannon divergence (up to a constant factor for $\alpha = \frac{1}{2}$) with Jeffreys divergence ($\alpha \rightarrow 0$).

Theorem 1 *There exists a parametric family of symmetric information-theoretic divergences $\{sKL^{(\alpha)}\}_\alpha$ that unifies both Jeffreys J -divergence ($\alpha \rightarrow 0$) with Jensen-Shannon divergence ($\alpha = \frac{1}{2}$).*

This result can be obtained by considering skew average of distributions instead of the one-half of Eq. 15:

$$L_\alpha(p : q) = \frac{H((1-\alpha)p + \alpha q) - H(p)}{\alpha(1-\alpha)} \geq 0 \quad (39)$$

Then it comes out that (see Eq. 7)

$$\text{sKL}^{(\alpha)}(p, q) = \frac{1}{2\alpha(1-\alpha)}(L_\alpha(p : q) + L_\alpha(q : p)). \quad (40)$$

Note that $L_{\frac{1}{2}}(p : q) = 4K(p : q)$. The scaling factor is due to historical convention. However L_α is in general not a ϕ -divergence.

An alternative description of the symmetric family is given by

$$S_F^{(\alpha)}(p, q) = \frac{2}{1-\alpha^2} \left(F(p) + F(q) - F\left(\frac{1-\alpha}{2}p + \frac{1+\alpha}{2}q\right) - F\left(\frac{1+\alpha}{2}p + \frac{1-\alpha}{2}q\right) \right). \quad (41)$$

It can be checked that $\text{sJ}_F^{(\alpha)}(p, q) = S_F^{(\alpha')}(p, q)$ for $\alpha' = 1 - 2\alpha$.

Many parametric distributions follow a regular structure called exponential families. In the next section, we shall link that class of symmetric sJ^α -divergences to equivalent symmetric α -Bhattacharyya divergences computed on the distribution parameter space.

3 Case of statistical exponential families

Many common statistical distributions are handled in the unified framework of exponential families [Nielsen and Nock(2009), Nielsen and Boltz(2011)]. A distribution is said to belong to an exponential family E_F , if its *parametric* density can be canonically rewritten as

$$p_F(x; \theta) = \exp(\langle t(x), \theta \rangle - F(\theta) + k(x)), \quad (42)$$

where θ describes the member of the exponential family $E_F = \{p_F(x; \theta) \mid \theta \in \Theta\}$, characterized by the log-normalizer $F(\theta)$, a convex differentiable function. $\langle x, y \rangle$ denotes the inner-product (e.g., $x^T y$ for vectors, see [Nielsen and Nock(2009), Nielsen and Boltz(2011)]). $t(x)$ is the sufficient statistic.

Discrete d -dimensional distributions (corresponding to frequency histograms with d non-empty bins in visual applications) are multinomials, an exponential

family with the dimension of the natural space Θ being $d - 1$ (the order of the family). In information retrieval [Rubner, Puzicha, Tomasi and Buhmann(2001)], one often needs to perform clustering on frequency histograms for building a codebook to perform efficiently retrieval queries (eg., bag-of-words method [Fei-Fei and Perona(2005)]).

It is known that the Kullback-Leibler divergence of members $p \sim E_F(\theta_p)$ and $q \sim E_F(\theta_q)$ of the same exponential family E_F is equivalent to a Bregman divergence on the swapped natural parameters [Banerjee et al.(2005)Banerjee, Merugu, Dhillon, and Ghosh]:

$$\text{KL}(p_F(x; \theta_p) : p_F(x; \theta_q)) = B_F(\theta_q : \theta_p), \quad (43)$$

where $B_F(\theta_q : \theta_p) = F(\theta_q) - F(\theta_p) - \langle \theta_q - \theta_p, \nabla F(\theta_p) \rangle$.

The Jeffreys J -divergence on members of the same exponential family (lhs.) can be computed as a symmetrized Bregman divergence, yielding an equivalent calculation on the natural parameter space (rhs.):

$$J(p_F(x; \theta_p), p_F(x; \theta_q)) = (\theta_p - \theta_q)^T (\nabla F(\theta_p) - \nabla F(\theta_q)) \quad (44)$$

Note that although the product of two exponential families is an exponential family, it is *not* the case for the mixture of two exponential families. Indeed, the mixture $(1 - \alpha)p + \alpha q$ does not in general belong to E_F . Therefore, the Jensen-Shannon divergence on members of the same exponential family *cannot* be computed directly from the natural parameters, since it requires to compute the entropy of the mixture distribution (with no known generic closed form):

$$\text{JS}(p = p_F(x; \theta_p), q = p_F(x; \theta_q)) = H\left(\frac{p + q}{2}\right) - \frac{H(p) + H(q)}{2}, \quad (45)$$

In fact, it turns out that Eq. 43 is the limit case of the property that α -skew Bhattacharyya divergence $B^{(\alpha)}$ of members $p = p_F(x; \theta_p)$ and $q = p_F(x; \theta_q)$ of the same exponential family E_F is equivalent to a α -Jensen divergence defined on the natural parameters [Nielsen and Boltz(2011)]:

$$B^{(\alpha)}(p_F(x; \theta_p) : p_F(x; \theta_q)) = -\log \int p_F(x; \theta_p)^\alpha p_F(x; \theta_q)^{1-\alpha} dx, \quad (46)$$

$$= J_F^{(\alpha)}(\theta_p : \theta_q), \quad (47)$$

with the α -Jensen divergence defined on the distribution d -dimensional parameter vectors as

$$J_F^{(\alpha)}(\theta_p : \theta_q) = \frac{1}{\alpha(1 - \alpha)} \sum_{i=1}^d ((1 - \alpha)F(\theta_p^i) + \alpha F(\theta_q^i) - F((1 - \alpha)\theta_p^i + \alpha\theta_q^i)). \quad (48)$$

We can therefore symmetrize α -skew Bhattacharyya divergences:

$$\begin{aligned} \text{sB}^{(\alpha)}(p_F(x; \theta_p), p_F(x; \theta_q)) &= \frac{1}{2}(B^{(\alpha)}(p_F(x; \theta_p) : p_F(x; \theta_q)) + B^{(\alpha)}(p_F(x; \theta_q) : p_F(x; \theta_p))) \\ &= -\frac{1}{2} \log \left(\int p^\alpha(x) q^{1-\alpha}(x) dx \right) \left(\int p^{1-\alpha}(x) q^\alpha(x) dx \right) \end{aligned} \quad (50)$$

$$= \alpha(1 - \alpha) \text{sJ}_F^{(\alpha)}(\theta_p, \theta_q), \quad (51)$$

and obtain equivalently a symmetrized skew Jensen divergence on the natural parameters.

Theorem 2 *The symmetrized skew α -Bhattacharyya divergence on members of the same exponential family is equivalent to a symmetrized skew α -Jensen divergence defined for the log-normalizer and computed in the natural parameter space.*

Let us now consider computing centroidal centers (say, for k -means clustering applications [Banerjee et al.(2005)Banerjee, Merugu, Dhillon, and Ghosh]).

4 Symmetrized skew α -Jensen centroids

Consider the discrete symmetrized α -Jensen divergences (not any more on distributions but on d -dimensional parameter vectors). In particular, we get for separable divergences:

$$\text{sJ}_F^{(\alpha)}(x, y) = \frac{1}{2\alpha(1-\alpha)} \sum_{i=1}^d \left(F(x^i) + F(y^i) - F(\alpha x^i + (1-\alpha)y^i) - F((1-\alpha)x^i + \alpha y^i) \right). \quad (52)$$

This family of discrete measures includes the *extended Kullback-Leibler divergence* for unnormalized distributions by setting $F(x) = x \log x$. The *barycenter* b of n points p_1, \dots, p_n is defined as the (unique) point that minimizes the weighted average distance:

$$b = \arg \min_c \sum_{i=1}^n w_i \times \text{sJ}_F^{(\alpha)}(p_i, c), \quad (53)$$

for $w = (w_1, \dots, w_n)$ a normalized weight vector ($\forall i, w_i > 0$ and $\sum_i w_i = 1$). In particular, choosing $w_i = \frac{1}{n}$ for all i yields by convention the *centroid*. Note that the multiplicative factor in the energy function of Eq. 53 does not impact

the minimum. Thus we need to equivalently minimize:

$$\min_c E(c) = \min_c \sum_{i=1}^n w_i (F(p_i) + F(c) - F(\alpha p_i + (1 - \alpha)c) - F(\alpha c + (1 - \alpha)p_i)). \quad (54)$$

Removing the constant terms (i.e., independent of c), this amounts to minimize the following energy functional ($\sum_i w_i = 1$):

$$\operatorname{argmin}_c E(c) = \operatorname{argmin}_c F(c) - \sum_i w_i (F(\alpha p_i + (1 - \alpha)c) + F(\alpha c + (1 - \alpha)p_i)). \quad (55)$$

Since F is convex, E is the minimization of a sum of a convex function plus a concave function. Therefore, we can apply the ConCave-Convex Procedure [Sriperumbudur and Lanckriet(2009)] (CCCP) that guarantees to converge to a minimum. We thus bypass using a gradient steepest descent numerical optimization that requires to tune a learning step parameter.

Initializing

$$c_0 = \sum_{i=1}^n w_i p_i \quad (56)$$

to the Euclidean barycenter, we iteratively update as follows:

$$\nabla F(c_{t+1}) = \sum_{i=1}^n w_i ((1 - \alpha) \nabla F(\alpha p_i + (1 - \alpha)c_t) + \alpha \nabla F(\alpha c_t + (1 - \alpha)p_i)), \quad (57)$$

That is,

$$c_{t+1} = (\nabla F)^{-1} \left(\sum_{i=1}^n w_i ((1 - \alpha) \nabla F(\alpha p_i + (1 - \alpha)c_t) + \alpha \nabla F(\alpha c_t + (1 - \alpha)p_i)) \right) \quad (58)$$

(Observe that since F is strictly convex, its Hessian $\nabla^2 F$ is positive-definite so that the reciprocal gradient is ∇F^{-1} is well-defined, see [Rockafeller(1969)].)

In the limit case, we get the following *fixed point* equation:

$$c^* = (\nabla F)^{-1} \left(\sum_{i=1}^n w_i ((1 - \alpha) \nabla F(\alpha p_i + (1 - \alpha)c^*) + \alpha \nabla F(\alpha c^* + (1 - \alpha)p_i)) \right). \quad (59)$$

This rule is a quasi-arithmetic mean, and can alternatively be initialized using $c'_0 = \nabla F^{-1}(\sum_{i=1}^n w_i \nabla F(p_i))$ instead of c_0 .

Let us instantiate this updating rule for $\alpha = \frac{1}{2}$ and $w_i = \frac{1}{n}$ on Shannon and Burg information functions:

Shannon information $F(x) = x \log x - x$	Burg information $F(x) = -\log x$
$\nabla F(x) = \log x, (\nabla F)^{-1}(x) = \exp x$	$\nabla F(x) = -1/x, (\nabla F)^{-1}(x) = -1/x$
$c_{t+1} = \left(\prod_{i=1}^n \frac{c_t + p_i}{2} \right)^{\frac{1}{n}}$	$c_{t+1} = \frac{n}{\sum_{i=1}^n \frac{2}{c_t + p_i}}$
\rightarrow Geometric update	\rightarrow Harmonic update

Note that for Jeffreys ($\alpha = 0$) and Jensen-Shannon ($\alpha = \frac{1}{2}$) divergences, the energy function is convex, and therefore the minimum is necessarily *unique*. (In fact, both Jeffreys and Jensen-Shannon are two instances of the class of convex Ali-Silvey-Csiszár divergences [Csiszár(1967), Ali and Silvey(1966)].)

Since α -JS divergences are ϕ -divergences (convex in both arguments), the barycenter with respect to α -JS is unique, and can be computed alternatively using any convex optimization technique. Ben-Tal et al. [Ben-Tal et al.(1989) Ben-Tal, Charnes, and Telmer] called those center points entropic means; They consider scalar values that can be extended to dimension-wise separable divergences, but *not* to normalized nor continuous distributions.

Theorem 3 *The centroid of members of the same exponential family with respect to the symmetrized α -Bhattacharyya divergence can be computed equivalently as the centroid of their natural parameters with respect to the symmetrized α -Jensen divergence using the concave-convex procedure.*

Note that for members of the same exponential family, both c_0 or c'_0 initializations are interpreted as left-sided or right-sided Kullback-Leibler centroids [Nielsen and Nock(2009)].

5 Experimental results

Statistical distances play important roles either in *supervised* classification tasks (e.g., interclass distance measure for feature subset selection methods [(Molina et al. (2002))Molina] or in *unsupervised* clustering (e.g., centroid-based k -means [Banerjee et al.(2005) Banerjee, Merugu, Dhillon, and Ghosh]). In unsupervised settings, statistical divergences are both used in the clustering preprocessing stage for building a codebook (e.g., using k -means algorithm [Banerjee et al.(2005) Banerjee, Merugu, Dhillon, and Ghosh]), and when answering on-line queries (e.g., classification using the nearest neighbor rule).




Fig. 2. Performance of the symmetrized α -skew Jensen divergences for a binary classification task (y -axis, percentage of correct classification) with respect to $\alpha \in [0, \frac{1}{2}]$ (x -axis).

Since we proposed a novel parametric family of statistical symmetric divergences linking *continuously* the Jeffreys divergence to the Jensen-Shannon divergence, let us study the impact of the α parameter in a toy application. Namely, we consider binary classification of images: That is, given a set of annotated images either with tag 1 or tag 2, perform classification of images using the nearest neighbor rule. We use the Caltech 101 database [Fei-Fei and Perona(2005)] that consists of 101 categories with about 40 to 800 images per category, and select the **airplanes** (800 images) and **Faces** (436 images) categories. For each color image, we choose the intensity histogram² as its feature vector. We compute a centroid for each of the **airplane**/**Faces** categories, and classify all images according to the nearest neighbor rule between those two class histogram centroids and the query image histogram with respect to the symmetrized α -skew Jensen divergence. Figure 2 displays the performance plot (correct percentage of classification) of this binary classification task. We empirically observe that performance may vary with α as expected, and that the best α needs to be tuned according to the training data hinting at the underlying geometry of data-sets. Here, the best correct classification rate (about 88%) is obtained for $\alpha = \frac{1}{4}$, that is for the mid-divergence between Jeffreys and Jensen-Shannon divergences.

6 Concluding remarks

In this paper, we have introduced a novel parametric family of symmetric divergences based on Jensen’s inequality called *symmetrized α -skew Jensen divergences*. Instantiating this family for the Shannon information generator, we

² We convert (R, G, B) colors into corresponding intensities $I = 0.3R + 0.596G + 0.11B$, and ensure (by adding a small non-zero constant) that histogram bins are never empty in order to have proper multinomial distributions.

have exhibited a one-parameter family of symmetrized Kullback-Leibler divergences. Furthermore, we showed that for distributions belonging to the *same* exponential family, the symmetrized α -Bhattacharyya divergence amounts to compute a symmetrized α -Jensen divergence defined on the parameter space, thus yielding a closed-form formula. We then reported an iterative algorithm for computing the centroid with respect to this class of divergences.

For applications like information retrieval requiring symmetric statistical distances, the choice is therefore not anymore to decide between Jeffreys or Jensen-Shannon divergences, but rather to choose or tune the best α parameter according to the application and input data.

Acknowledgments

The author would like to thank the anonymous reviewer for his/her valuable comments and suggestions.

References

- [Ali and Silvey(1966)] Ali, S. M., Silvey, S. D., 1966. A general class of coefficients of divergence of one distribution from another. *Journal of the Royal Statistical Society, Series B* 28, 131–142.
- [Banerjee et al.(2005)] Banerjee, Merugu, Dhillon, and Ghosh]
Banerjee, A., Merugu, S., Dhillon, I. S., Ghosh, J., 2005. Clustering with Bregman divergences. *J. Mach. Learn. Res.* 6, 1705–1749.
- [Ben-Tal et al.(1989)] Ben-Tal, Charnes, and Teboulle] Ben-Tal, A., Charnes, A., Teboulle, M., 1989. Entropic means. *Journal of Mathematical Analysis and Applications* 139 (2), 537 – 551.
- [Cover and Thomas(1991)] Cover, T. M., Thomas, J. A., 1991. *Elements of information theory*. Wiley-Interscience, New York, NY, USA.
- [Csiszár(1967)] Csiszár, I., 1967. Information-type measures of difference of probability distributions and indirect observation. *Studia Scientiarum Mathematicarum Hungarica* 2, 229–318.
- [Deselaers, Keysers, and Ney(2008)] Deselaers, T., Keysers, D. and Ney, H. 2008. Features for image retrieval: an experimental comparison. *Inf. Retr.* 11(2), pp. 77–107.
- [Fei-Fei and Perona(2005)] Fei-Fei, L., Perona, P., June 2005. A Bayesian hierarchical model for learning natural scene categories. Vol. 2. *IEEE Computer Society*, pp. 524–531.

- [Jeffreys(1973)] Jeffreys, H., 1973. Scientific Inference. Cambridge University Press.
- [Jensen(1906)] Jensen, J. L. W. V., December 1906. Sur les fonctions convexes et les inégalités entre les valeurs moyennes. *Acta Mathematica* 30 (1), 175–193.
- [Kullback and Leibler(1951)] Kullback, S., Leibler, R. A., 1951. On information and sufficiency. *Annals of Mathematical Statistics* 22, 49–86.
- [Lin(1991)] Lin, J., 1991. Divergence measures based on the Shannon entropy. *IEEE Transactions on Information Theory* 37, 145–151.
- [(Liu et al. (2010))Liu, Vemuri, Amari and Nielsen] Liu, M. and Vemuri, B. C. and Amari, S.-i and Nielsen F., 2010. Total Bregman Divergence and its Applications to Shape Retrieval, *IEEE CVPR* 2010.
- [(Molina et al. (2002))Molina, Belanche and Nebot] Molina, L. C., Belanche, L. and Nebot, A. 2002. Feature Selection Algorithms: A Survey and Experimental Evaluation. In *Proc. IEEE International Conference on Data Mining*, pp. 306–313.
- [Nielsen and Boltz(2011)] Nielsen, F., Boltz, S., 2011. The Burbea-Rao and Bhattacharyya centroids. *IEEE Transactions on Information Theory* 57(8), pp. 5455–5466.
- [Nielsen and Nock(2009)] Nielsen, F., Nock, R., 2009. Sided and symmetrized Bregman centroids. *IEEE Transactions on Information Theory* 55(6), pp. 2048–2059.
- [Rényi(1961)] Rényi, A., 1961. On measures of entropy and information. In *Proc. 4th Berkeley Symp. Math. Stat. and Prob.* Vol. 1. pp. 547–561.
- [Rockafeller(1969)] Rockafeller, R. T., *Convex analysis*, Princeton University Press, 1969.
- [Rubner, Puzicha, Tomasi and Buhmann(2001)] Y. Rubner, J. Puzicha, C. Tomasi, and M. Buhmann. 2001. Empirical evaluation of dissimilarity measures for color and texture. *Comput. Vis. Image Underst.* 84(1), pp. 25–43.
- [Sriperumbudur and Lanckriet(2009)] Sriperumbudur, B., Lanckriet, G., 2009. On the convergence of the concave-convex procedure. In: Bengio, Y., Schuurmans, D., Lafferty, J., Williams, C. K. I., Culotta, A. (Eds.), *Advances in Neural Information Processing Systems* 22. pp. 1759–1767.